arXiv:1611.08927v1 [cs.GT] 27 Nov 2016

The Complexity of Succinct Elections

Zack Fitzsimmons
College of Computing and Inf. Sciences
Rochester Inst. of Technology
Rochester, NY 14623

Edith Hemaspaandra
Dept. of Computer Science
Rochester Inst. of Technology
Rochester, NY 14623

November 27, 2016

Abstract

The computational study of elections generally assumes that the preferences of the electorate come in as a list of votes. Depending on the context, it may be much more natural to represent the list succinctly, as the distinct votes of the electorate and their counts. We consider how this succinct representation of the voters affects the computational complexity of election problems. Though the succinct representation may be exponentially smaller than the nonsuccinct representation, we find only one natural case where the complexity increases, namely the complexity of winner determination for Kemeny elections. This is in sharp contrast to the case where each voter has a weight, where the complexity usually increases.

1 Introduction

Elections are an important and widely used tool for determining an outcome given the preferences of several agents. Unfortunately, an important negative result from social-choice theory, the Gibbard-Satterthwaite Theorem states that every reasonable election system is manipulable [Gib73, Sat75]. However, even though every reasonable election system is manipulable, it may be computationally infeasible to determine if a manipulation is possible. Bartholdi, Tovey, and Trick [BTT89] started the computational study of the manipulation problem and later introduced the study of a family of manipulative attacks, denoted control, which models the actions of an election chair with control over the structure of the election and wants to ensure his or her preferred outcome [BTT92].

In most of the computational studies on elections, the preferences of the voters are represented as a list of votes. Though this may be a reasonable representation for paper ballots in political elections, in artificial intelligence applications a more succinct representation where the preferences of the electorate are represented as a list of distinct votes and their counts may be more natural. For example, this representation is used by the online preference repository PrefLib for election data [MW13].

We consider how this succinct representation of the votes can affect the complexity of different election problems, and contrast this with the case of weighted voters.

Though the succinct representation may be exponentially smaller than the nonsuccinct representation, we find that in surprisingly few cases the complexity increases. Related work that considers succinct votes did not find a case where the complexity increases [FHHR09, FHH09, HHR09, FHHR11]. We explain this phenomenon by showing that many common proof techniques that show that election problems are in P can be adapted to the succinct case (see Section 3).

However, we also show that it is possible for the complexity to increase. For example, consider the following system X. Every voter votes by an approval vector (i.e., states an approval or disapproval for each candidate). A candidate wins if they are approved by exactly half of the voters. Though simple, this election system is not particularly natural. It has been chosen to show the increase in complexity: Consider the problem of constructive control by adding candidates (CCAC) for election system X. Given an election where the votes are represented succinctly, a set of unregistered candidates, and a preferred candidate p it is NP-complete to ensure that p wins the election using system X by adding unregistered candidates to the election. This is because this problem is basically subset sum, which is in P for unary numbers by dynamic programming and is NP-complete for binary numbers [GJ79]. (See Theorem 19 for a straightforward proof of this example.)

We found a natural case where the complexity increases. The complexity of determining the winner in a Kemeny election, which is well-known to be Θ_2^p-complete for the case of nonsuccinct votes [HSV05], is Δ_2^p-complete for succinct votes. This solves an open problem from the previous work by Hemaspaandra et al. [HSV05]. (See Section 4.)

2 Preliminaries

A *(nonsuccinct) election* is defined as a finite set of candidates C, and a list of voters V, where each voter $v \in V$ has a preference order over the candidates. In a *succinct election* V is not a list of voters, but instead a list of distinct votes v (preference orders) and their positive integer count $\kappa(v)$. In a *weighted election* V is still a list of voters, but now each $v \in V$ has an associated positive integer weight $\omega(v)$, and can be thought of as a coalition of $\omega(v)$ voters all voting the same. It is important to note that weights are indivisible, while the counts in a succinct election are not. We assume that the preference order of each voter is a total order, i.e., he or she strictly ranks each candidate from most to least preferred.

An election system $\mathcal{E}$ maps an election to set of winners, where the winners can be any subset of the candidate set. The problem $\mathcal{E}$-Winner is defined below.

Name: $\mathcal{E}$-Winner

Given: An election (C, V) and a candidate $p \in C$.

Question: Is p a winner of (C, V) using election system $\mathcal{E}$?

In the succinct case, $\mathcal{E}$-Succinct-Winner, V is represented succinctly, and in the weighted case, $\mathcal{E}$-Weighted-Winner, each voter has a corresponding positive integer weight.

A scoring vector $\langle \alpha_1, \alpha_2, \ldots, \alpha_m \rangle$, $\alpha_i \geq \alpha_{i+1}$ defines an election system over m candidates. Each candidate receives α_i points for each vote where they are ranked ith, and the candidate(s) with the highest score win.

A *pure scoring rule* defines a family of scoring vectors where the m-candidate scoring vector can be computed in polynomial time in m, and the $m+1$-candidate scoring vector can be obtained from the m-candidate scoring vector by adding a single coefficient [BD09].

2.1 Manipulative Actions

We examine the complexity of the following manipulative actions on elections. For each problem, we present the definition for the nonsuccinct case, and describe how the succinct and weighted cases differ.

Manipulation is the most widely studied manipulative action on elections. The computational study of manipulation was introduced by Bartholdi, Tovey, and Trick [BTT89], and was later extended for the coalitional, destructive, and weighted cases by Conitzer, Sandholm, and Lang [CSL07]. We present the definition of constructive unweighted coalitional manipulation (CUCM) below.

Name: $\mathcal{E}$-CUCM

Given: A set of candidates C, a list of nonmanipulative voters V_1, a list of k manipulators V_2,[1] and a preferred candidate $p \in C$.

Question: Is there a way to set the preferences of the manipulators so that p is a winner of the election $(C, V_1 \cup V_2)$ using election system $\mathcal{E}$?

The weighted case (CWCM) is essentially the same definition as above, except each of the voters (both the nonmanipulators and the manipulators) have an associated positive integer weight. In the corresponding succinct case, the nonmanipulator votes are represented succinctly, and the manipulators are represented as k encoded in binary.

Electoral control denotes the family of manipulative actions that consider an agent with control over the structure of the election, called the election chair, who wants to ensure his or her preferred outcome [BTT92]. Below we define constructive control by adding voters (CCAV), which is a very natural case of control and can be thought of as modeling get-out-the-vote drives.

Name: $\mathcal{E}$-CCAV

Given: A set of candidates C, a list of registered voters V, a list of unregistered voters U, an add limit $k \in \mathbb{N}$, and a preferred candidate $p \in C$.

Question: Does there exist a list U' of unregistered voters $U' \subseteq U$ such that U' consists of at most k unregistered voters and p is a winner of $(C, V \cup U')$ using election system $\mathcal{E}$?

[1] This can also be thought of as specifying the number of manipulators, k, in unary.

In the standard model of weighted voter control the parameter k denotes the number of weighted voters the chair can add/delete [FHH15]. In the succinct case, the only change from the $\mathcal{E}$-CCAV definition above is that the voters (registered and unregistered) are represented succinctly.

2.2 Computational Complexity

We assume that the reader is familiar with the complexity classes P and NP. We also have several results concerning the complexity classes Θ_2^p and Δ_2^p. Δ_2^p denotes the class of problems solvable in P with access to an NP oracle. Θ_2^p is a subset of Δ_2^p, where the P-machine can ask one round of parallel queries to its NP oracle. Θ_2^p is equivalent to $\mathrm{P}^{\mathrm{NP[log]}}$, the class of problems solvable in P with only a log number of queries to an NP oracle [Hem89].

3 Adapting Approaches

The general theme of this paper is that in surprisingly few cases, the complexity of election problems increases when we allow succinct representation. In this section we will discuss several common approaches used to show that election problems are in P (for the nonsuccinct case), and describe how they can be adapted for the case where voters are represented succinctly, sometimes straightforwardly and sometimes in a more complicated way.

In particular, we show how greedy approaches, limited brute-forcing, network flow, and edge matching/cover techniques can be adapted. To showcase these adaptations, we will show that the dichotomy result for constructive control by adding voters (CCAV) for pure scoring rules [HHS14] holds for the succinct case. Assuming P $\neq$ NP, the following pure scoring rules are asymptotically the only cases where CCAV is in P, both for the nonsuccinct and the succinct case.

- $\langle \alpha, \beta, 0, \ldots, 0 \rangle$, where $\alpha > \beta$.
- t-Approval, $\langle 1^t, 0, \ldots, 0 \rangle$, where $t \leq 3$.
- t-Veto, $\langle 1, \ldots, 1, 0^t \rangle$, where $t \leq 2$.
- $\langle 2, 1, \ldots, 1, 0 \rangle$.

In contrast, all weighted cases are NP-complete, except triviality ($\langle 0, \ldots, 0 \rangle$), 1-approval (plurality), 2-approval, and 1-veto [FHH15, Lin12].

The following case demonstrates the adaptation of greedy algorithms as well as how to handle limited brute-forcing.

Theorem 1 *For $\alpha > \beta \geq 0$, $\langle \alpha, \beta, 0, \ldots, 0 \rangle$-CCAV is in* P *for succinct elections.*

Proof. We follow the proof from [HHS14] and show how to adapt this to the succinct case.

The proof shows that there is a constant ℓ such that it is never better to add ℓ voters that put p second than it is to add ℓ pairwise different voters that put p first (assuming that voters that rank the first two candidates identically also rank all other candidates identically). Let V_1 be the set of voters that put p first and let V_2 be the set of voters that put p second.

So, if p can be made a winner, p can be made a winner by adding at most ℓ voters from V_2 or in a way that does not leave ℓ pairwise different voters in V_1 unused.

In the first case, we can brute-force over all sets of V_2 voters of size $\leq \ell$. Add these voters. We are then left with the problem of checking whether we can add at most k' V_1 voters to an election to make p a winner, where k' is k minus the number of V_2 voters added. In Hemaspaandra et al. [HHS14], this is done by brute-forcing over every j in $\{0, \ldots, k'\}$ and checking whether we can add j voters from V_1 to make p a winner. If we know j, we know what the score of p after control will be, and so the last part can then be done greedily, by for each candidate adding as many voters that put a second as possible one voter at the time (until we have added j voters total).

We have two problems in the succinct case. The first is that we cannot in polynomial time brute-force over all possible values of j, since k' is not polynomially bounded. The second problem is that we don't have the time to add voters one at the time. The second problem can of course easily be fixed by, for each candidate a, adding as many voters with a in second place as possible all at once: If s_a is the initial score of a and fs_p is the final score of p (note that $fs_p = s_p + j\alpha$), we can add at most $\lfloor (fs_p - s_a)/\beta \rfloor$ voters with a second.

To handle the first problem, note that we can formulate the problem as an integer linear program with one variable (j), namely, there is a j, $0 \leq j \leq k'$ such that

$$\sum_{a \in C - \{p\}} \min(\lfloor (s_p + j\alpha - s_a)/\beta \rfloor, v_a) \geq j,$$

where v_a is the number of voters in V_1 that rank a second.[2] This integer linear program can be solved in polynomial time by [Len83].

It remains to show that the second case, where p can be made a winner in a way that does not leave ℓ pairwise different voters in V_1 unused can be determined in polynomial time in the succinct case.

For the nonsuccinct case, in Hemspaandra et al. [HHS14], this case is handled by brute-forcing over all sets S of at most $\ell - 1$ candidates who are ranked second by unused V_1 voters, and then brute-forcing over all possibilities for these candidates, i.e., by brute-forcing over all functions $u : S \to \{1, \ldots, \|V_1\|\}$. This way, we loop over all cases of adding V_1 voters. After adding the V_1 voters, we need to see if p can be made a winner by adding V_2 voters. This can be done in a similar way as adding the V_1 voters in the first case above.

[2] Notice that the ILP above checks that if we add the maximum number of voters such that the score of each of the non-p candidates is at most $s_p + j\alpha$, then we have added at least j voters. This implies that we can add j voters in such a way that p is still a winner.

In the succinct case, we can still loop over all sets S of at most $\ell - 1$ candidates. We can not brute-force all functions u, but this will turn again into an integer linear program, this time with $\ell - 1$ variables (for the u values) and one more variable to handle the V_2 voters similarly to how we handled the V_1 voters in the first case above. This integer linear program can be solved in polynomial time by [Len83]. ❑

Theorem 2 *CCAV is in* P *for t-approval when $t \leq 3$ and for t-veto when $t \leq 2$ for succinct elections.*

Proof. The results for 1-approval, 2-approval, and 1-veto follow via simple greedy algorithms that can easily be adapted to work in the succinct case.

3-approval-CCAV was shown to be in P by Lin [Lin12] using a reduction to Simple b-Edge Matching for Multigraphs. The essence of the reduction is that every unregistered voter of the form $\{p, a, b\} > \cdots$ (by which we mean a voter who gives a point to p, a, and b) corresponds to an edge (a, b) in the constructed multigraph. For every candidate $c \neq p$, c is a vertex in the graph and $b(c)$ is the final score of p minus the initial score of c (i.e., the number of points that we can add to c while keeping c's score less than or equal to p's score). In the succinct case, we would have too many edges in the graph. However, Capacitated b-Edge Matching (for graphs where edges have integer capacities) is also in P. So, we simply set the capacity of edge (a, b) equal to the number of unregistered voters of the form $\{p, a, b\} > \cdots$, and we let $b(c)$ be the final score of p minus the initial score of c as previously.

Similarly, 2-veto-CCAV was shown to be in P by reduction to Simple b-Edge Cover for Multigraphs. In that reduction, every unregistered voter $\cdots > a > b$ corresponds to an edge (a, b) in the constructed multigraph. Again, we can modify this construction to a reduction to Capacitated b-Edge Cover, which is also in P, by letting the capacity of edge (a, b) be the number of unregistered voters $\cdots > a > b$. ❑

Another common technique to show that election problems are in P is network flow.

Theorem 3 *$\langle 2, 1, \ldots, 1, 0\rangle$-CCAV is in* P *for succinct elections.*

Proof. The essence of the nonsuccinct proof is to (after some easy preprocessing) build a min-cost flow network. The capacities of the edges are computed from the scores of the candidates and the multiplicities of the voters and are polynomially bounded in the size of the input. In the succinct case, we can use the exact same network. The capacities are now binary integers, but min-cost network flow is still in P in that case. ❑

3.1 Dynamic Programming

There is one other common technique that is used to prove that election problems are in P. This is dynamic programming. As alluded to in the introduction, dynamic programming

approaches do not generalize to the succinct case. This does not mean that the succinct case is not in P. For example, Hemaspaandra and Schnoor [HS16] show that for all pure scoring rules with a constant number of different coefficients, manipulation is in P. This is shown by dynamic programming. And this algorithm is not in P for the succinct case, not even for restricted versions of this problem.

Of course, as just stated, we may not need dynamic programming. For example, it is immediate that plurality and veto can be easily solved in polynomial time, in the nonsuccinct and the succinct case. For a more involved example, for scoring rule $\langle 0, \ldots, 0, -1, -2\rangle$, one can greedily put a candidate $c \neq p$ with highest score in a -2 spot, and after filling all -2 spots, in a -1 spot. Note that this simple greedy algorithm does not directly generalize to the succinct case, since we do not have enough time to fill slots one-at-a-time.

Still, in this particular case it is not hard to see that the succinct case is in P: Greedily put a highest scoring candidate $c \neq p$ in $\lfloor(\text{score}(c) - \text{score}(p))/2\rfloor$ -2 spots (or in as many of those as available). After filling all -2 spots, fill the -1 spots.

Things quickly get more complicated though. And there is a clear limit to how far this will generalize, since Borda manipulation is NP-complete.

Still, we conjecture that all these manipulation cases are in P for succinct elections.

Again, this is in contrast to the weighted case, which is in P only for triviality and plurality. This holds even for fixed numbers of candidates, as explained in the next section.

3.2 Fixed Number of Candidates

It is reasonable to assume that the number of candidates in an election may be fixed. For the case of weighted voters many problems are hard, even when the number of candidates is fixed.

Theorem 4 *1. m-candidate α-CWCM is* NP*-complete for every scoring rule α that is not plurality or triviality [HH07].*

2. m-candidate α-$CC\left[\begin{smallmatrix}\text{A}\\ \text{D}\end{smallmatrix}\right]$ V is NP*-complete for every scoring rule $\alpha = \langle\alpha_1, \ldots, \alpha_m\rangle$, when $\|\alpha_1, \ldots, \alpha_m\| \geq 3$, for weighted elections [FHH15].*

Faliszewski et al. [FHH09] showed that for succinct votes that manipulation is in P for any fixed number of candidates by describing an integer linear program with a fixed number of variables.

Theorem 5 *[FHH09] m-candidate α-CUCM is in* P *for every scoring rule α, for succinct elections.*

A similar approach can be used to show that for succinct votes, constructive control by adding/deleting voters is in P for every scoring rule for a fixed number of candidates. Details can be found in the appendix.

Theorem 6 *m-candidate α-$CC\left[\begin{smallmatrix}\text{A}\\ \text{D}\end{smallmatrix}\right]$ V is in* P *for every scoring rule $\alpha = \langle\alpha_1, \ldots, \alpha_m\rangle$, for succinct elections.*

4 Kemeny Elections

Kemeny elections were introduced in [Kem59]. It is the unique election system that is neutral, consistent, and Condorcet [YL78]. A candidate p is a Kemeny winner if p is ranked first in a Kemeny consensus. A Kemeny consensus is a linear order $>$ over the candidates that minimizes the Kendall's tau distance to V, i.e., that minimizes

$$\sum_{a,b \in C, a>b} \|\{v \in V \mid b >_v a\}\|,$$

where $>_v$ is the preference order of voter v.

Observation 7 *Kemeny-Weighted-Winner is equivalent to Kemeny-Succinct-Winner.*

Notice that the above observation holds since we do not modify the votes to score the election. This is not the case for Dodgson and Young elections, which we will discuss later in the section.

The Kemeny score of a candidate c is the minimum value of the Kendall's tau distance of all linear orders over the candidates that rank c first. And so p is a Kemeny winner if and only of p has a lowest Kemeny score.

Name: Kemeny-Score

Given: An election (C, V), a candidate $c \in C$, and an integer k.

Question: Is the Kemeny score of c at most k?

Kemeny-Score was shown to be NP-complete and Kemeny-Winner was shown NP-hard by a reduction from feedback arc set in [BTT89]. Kemeny-Winner was shown Θ_2^p-complete in [HSV05]. Membership in Θ_2^p is easy: Simply compute the Kemeny scores of all candidates with one round of parallel queries to Kemeny-Score and check that p's score is the lowest. Kemeny-Winner was shown to be Θ_2^p-complete by a chain of three reductions, ultimately showing that the Θ_2^p-complete problem Min-Card-Vertex-Cover-Compare reduces to Kemeny-Winner.

Theorem 8 *[HSV05] Kemeny-Winner is Θ_2^p-complete.*

Footnote 3 in [HSV05] points out that Kemeny-Succinct-Winner is in Δ_2^p and explicitly leaves the exact complexity of this problem as an open question. It is easy to see that Kemeny-Succinct-Winner is in Δ_2^p: Simply compute the Kemeny scores of all candidates using binary search on queries to Kemeny-Weighted-Score (which is in NP) and check that p's score is the lowest). We will now show that Kemeny-Succinct-Winner (and Kemeny-Weighted-Winner) are in fact Δ_2^p-complete.

We first define the following weighted version of Min-Card-Vertex-Cover-Compare.

Name: Weighted-Min-Card-Vertex-Cover-Compare

Given: Vertex-weighted graphs G and H such that $\omega(G) = \omega(H)$, where $\omega(G)$ denotes the weight of the graph.

Question: Is the weight of G's minimum-weight vertex cover less than or equal to the weight of H's minimum-weight vertex cover?

Lemma 9 *Weighted-Min-Card-Vertex-Cover-Compare polynomial-time many-one reduces to Kemeny-Succinct-Winner.*

Proof. This follows by careful inspection and modification of the proof from [HSV05]. In particular, since that proof consists of a chain of three reductions between Θ_2^p-complete problems, we need to define suitable Δ_2^p-complete weighted versions of the two intermediate problems and show that the weighted versions of the reductions still hold. In essence, we "lift" the constructions and proofs from Θ_2^p to Δ_2^p. This works surprisingly smoothly. Details can be found in the appendix. ❑

It remains to show the following lemma.

Lemma 10 *Weighted-Min-Card-Vertex-Cover-Compare is Δ_2^p-complete.*

Proof. Membership in Δ_2^p is easy to see: Use binary search to compute the the weight of G's and H's minimum-weight vertex covers. Hardness is much more complicated, since there are not many Δ_2^p-complete problems to choose from (significantly less than even Θ_2^p-complete problems). The closest available problem is MAXSATASG$_=$ [Wag87], which consists of pairs of 3cnf formulas with the same maximal satisfying assignment. Though Wagner does not specify what happens when the formulas are not satisfiable, it is easy to see from the proof of Δ_2^p-hardness of MAXSATASG$_=$ that we need look only at satisfiable formulas.

By negating the variables in the formulas, we obtain the following Δ_2^p-hard promise problem.[3]

Name: MINSATASG$_=$

Given: Two satisfiable 3cnf formulas $\phi(x_1, \ldots, x_n)$ and $\psi(x_1, \ldots, x_n)$.

Question: Do ϕ and ψ have the same lexicographically minimal satisfying assignment (viewed as an n-bit string)?

It is convenient to introduce the following intermediate Δ_2^p-complete problem.

Name: Weighted-Min-Card-Vertex-Cover-Equality

[3] A similar problem was used to show the only previously known Δ_2^p-completeness result in computational social choice, namely of Online-Veto-Weighted-Coalitional-Manipulation [HHR14].

Given: Vertex-weighted graphs G and H such that $\omega(G) = \omega(H)$, where $\omega(G)$ denotes the weight of the graph.

Question: Is the weight of G's minimum-weight vertex cover equal to the weight of H's minimum-weight vertex cover?

We will show that MINSATASG$_=$ reduces to Weighted-Min-Card-Vertex-Cover-Equality and that Weighted-Min-Card-Vertex-Cover-Equality reduces to Weighted-Min-Card-Vertex-Cover-Compare, which proves Lemma 10.

Let $\phi(x_1, \ldots, x_n)$ and $\psi(x_1, \ldots, x_n)$ be two satisfiable 3cnf formulas. Without loss of generality, assume that ϕ and ψ have the same number m of clauses (simply pad).

Let $f(\phi)$ be the graph computed by the standard reduction from 3SAT to Vertex-Cover from [Kar72]. Then $f(\phi)$ consists of $2n + 3m$ vertices: $\{x_i, \overline{x_i} \mid 1 \leq i \leq n\} \cup \{a_i, b_i, c_i \mid 1 \leq i \leq m\}$, and the following edges: $\{x_i, \overline{x_i}\}$ for $1 \leq i \leq n$ and if the jth clause in ϕ is $\ell_1 \vee \ell_2 \vee \ell_3$, where $\ell_r \in \{x_i, \overline{x_i} \mid 1 \leq i \leq n\}$, then we have edges $\{a_j, \ell_1\}$, $\{b_j, \ell_2\}$, $\{c_j, \ell_3\}$.

The properties of f that we need here, which follow immediately from the proof of [Kar72], are as follows.

1. $f(\phi)$ does not have a vertex cover of size less than $n + 2m$.

2. If W is a vertex cover of size $n + 2m$, then $W \cap \{x_i, \overline{x_i} \mid 1 \leq i \leq n\}$ corresponds to a satisfying assignment of ϕ, in the sense that $\|W \cap \{x_i, \overline{x_i}\}\| = 1$ for $1 \leq i \leq n$, and $\alpha_1 \cdots \alpha_n$ defined as $\alpha_i = 1$ if and only if $x_i \in W$ is a satisfying assignment for ϕ.

3. If α is a satisfying assignment for ϕ, then there is a vertex cover of size $n + 2m$ such that $W \cap \{x_i, \overline{x_i} \mid 1 \leq i \leq n\}$ corresponds to this assignment, i.e., the set $\{x_i \mid \alpha_i = 1\} \cup \{\overline{x_i} \mid \alpha_i = 0\}$ can be extended to a vertex cover of size $n + 2m$ by adding $2m$ vertices from $\{a_i, b_i, c_i \mid 1 \leq i \leq m\}$.

Now set the weights of the vertices as follows: $\omega(a_i) = \omega(b_i) = \omega(c_i) = 2^n$, $\omega(x_i) = 2^n + 2^{i-1}$, and $\omega(\overline{x_i}) = 2^n$. Note that for W a set of vertices, $\|W\| = \lfloor \omega(W)/2^n \rfloor$. In particular, a minimum-weight vertex cover will also have minimum size.

We will now show that n-bit string α is the smallest satisfying assignment of ϕ if and only if $f(\phi)$ has a minimum-weight vertex cover of weight $(n + 2m)2^n + \alpha$. (We interpret n-bit string α as a binary number between 0 and $2^n - 1$.)

If α is a satisfying assignment for ϕ, then by property 3 above the set $\{x_i \mid \alpha_i = 1\} \cup \{\overline{x_i} \mid \alpha_i = 0\}$ can be extended to a vertex cover of size $n + 2m$ by adding $2m$ vertices from $\{a_i, b_i, c_i \mid 1 \leq i \leq m\}$. The weight of this vertex cover is $(n + 2m)2^n + \alpha$. Since the minimum size of a vertex cover is $n + 2m$, the weight of a minimum-weight vertex cover is $(n + 2m)2^n + \beta$, for some β such that $0 \leq \beta < 2^n$. Such a β corresponds to a satisfying assignment by property 2 above.

Since ψ is also a satisfiable 3cnf formula over $x_1, \ldots, x_n$ with m clauses, it also holds that n-bit string α is the smallest satisfying assignment of ψ if and only if $f(\psi)$ has a minimum weight vertex cover of weight $(n + 2m)2^n + \alpha$.

This then implies that ϕ and ψ have the same minimal satisfying assignment if and only if $f(\phi)$ and $f(\psi)$'s minimum-weight vertex covers have the same weight. Note that $\omega(f(\phi)) = \omega(f(\psi))$. This completes the reduction from MINSATASG$_=$ to Weighted-Min-Card-Vertex-Cover-Equality.

It remains to show that Weighted-Min-Card-Vertex-Cover-Equality reduces to Weighted-Min-Card-Vertex-Cover-Compare.

This is straightforward. Given two vertex-weighted graphs G and H such that $\omega(G) = \omega(H)$, map this to the following two vertex-weighted graphs ($+$ denotes the disjoint union and $\times$ denotes the join of two graphs).

1. $\widehat{G} = X + G + H$, where X is a 2-vertex, 1-edge graph, where both vertices have weight $2\omega(G)$.
2. $\widehat{H} = (G \times H) + (G \times H)$.

Note that $\omega(\widehat{G}) = \omega(\widehat{H})$. Let mvc be the minimum vertex cover weight of a vertex-weighted graph. Then $mvc(\widehat{G}) = 2\omega(G) + mvc(G) + mvc(H)$. $mvc(\widehat{H}) = 2\min(\omega(G) + mvc(H), mvc(G) + \omega(H)) = 2\omega(G) + 2\min(mvc(G), mvc(H))$. It follows that $mvc(G) = mvc(H)$ if and only if $mvc(\widehat{G}) \leq mvc(\widehat{H})$. ❑

This completes the proof of the main theorem of this section.

Theorem 11 *Kemeny-Weighted-Winner and Kemeny-Succinct-Winner are Δ_2^p-complete.*

[DKNS01] shows that Kemeny-Winner is already NP-hard if we have four voters. In fact, one can easily combine the techniques from [HSV05] and [DKNS01] to obtain the following theorem. Details can be found in the appendix.

Theorem 12 *Kemeny-Winner for four voters is Θ_2^p-complete.*[4]

Do we get the same complexity jump if we look at the succinct and weighted case for four votes? Note that for the succinct case we allow many voters with the same vote, and so the Kemeny scores can be large. However, even though the scores can be large, it is easy to see that for each instance there are only a polynomial number of possible scores for each candidate, namely, if the multiplicities of the four votes are k_1, k_2, k_3, and k_4, the only possible scores are $\ell_1 k_1 + \ell_2 k_2 + \ell_3 k_3 + \ell_4 k_4$, where $0 \leq \ell_i \leq \|C\|(\|C\| - 1)/2$. For each candidate, simply query the Kemeny-Weighted-Score oracle for all these possible scores in parallel, and then determine whether p has the highest score.

[4]The definition of Kemeny winner can naturally be extended to voters voting any binary preference relation, and we note that the same argument as for Theorem 12 also shows the following. Details can be found in the appendix.

Theorem 13 *Kemeny-Winner for two voters who vote a partial order and Kemeny-Winner for one voter who votes a binary relation are Θ_2^p-complete.*

Theorem 14 *Kemeny-Weighted-Winner and Kemeny-Succinct-Winner for four votes are Θ_2^p-complete.*

We end this section by looking at two other voting systems whose winner problem is Θ_2^p-complete [HHR97, RSV03] namely, Dodgson voting [Dod76] and Young voting [You77].

The Young score of a candidate c is defined as the minimum number of voters that need to be deleted in order to make c a weak Condorcet winner. The Dodgson score of a candidate c is the minimum number of switches between adjacent candidates in voters such that c becomes a Condorcet winner. The candidates with the lowest scores are the winners. Unlike Kemeny, these winner problems depend on a modification of the votes, and so the succinct case and the weighted case do not coincide. In addition, one has to carefully think about what the weighted winner problem means: Do we charge by the voter (in analogy to the definition of weighted control) or do we charge by voter weight.

Clearly, by brute-force, we have the following.

Theorem 15 *For a fixed number of voters, Young-Weighted-Winner and Dodgson-Weighted-Winner are in P, both in the model where we charge by the voter and in the model where we charge by the vote-weight.*

The succinct case for Young winner is also in P, but this seems to need much more powerful machinery. So, though this is not a complexity jump from tractable to intractable, it seems to be a jump in the complexity of the algorithm.

Theorem 16 *For a fixed number of votes, Young-Succinct-Winner is in P.*

Proof. Let $v_1, \ldots, v_\ell$ be the different votes. Then the Young score of c is at most k if and only if there exist $k_1, \ldots, k_\ell$ such that $k_i \leq \kappa(v_i)$ and c is a weak Condorcet winner in the election system consisting of $\kappa(v_i) - k_i$ votes v_i. All this can be formulated as an integer linear program with ℓ variables, which can be solved in polynomial time by [Len83]. ❑

For the general weighted cases where we charge by the voter, it is easy to see that the scores are polynomially bounded, and so we get.

Theorem 17 *Young-Weighted-Winner and Dodgson-Weighted-Winner in the model where we charge by the voter are Θ_2^p-complete.*

It is easy to see that Young-Weighted-Winner and Dodgson-Weighted-Winner in the model where we charge by vote-weight as well as Young-Succinct-Winner and Dodgson-Succinct-Winner are in Δ_2^p, by using binary search to find the scores of all candidates.

We conjecture that, as in the case for Kemeny, we can "lift" the Θ_2^p-hardness proof for Dodgson and Kemeny winner to a Δ_2^p-hardness proof for the weighted case.

Conjecture 18 *Young-Weighted-Winner and Dodgson-Weighted-Winner in the model where we charge by vote-weight are Δ_2^p-complete.*

We do not think that this "lifting" will work for the succinct case and leave the exact complexity of Young-Succinct-Winner and Dodgson-Succinct-Winner, which we know to be Θ_2^p-hard and in Δ_2^p, as an open question.

5 Conclusions and Future Work

Overall, we found that when we allow succinct representation of the voters that the complexity of election problems rarely increases.

There are several interesting directions for future work. In the cases we found where the complexity increases when moving to a succinct representation, the weighted and succinct problems were equivalent. Will this always be the case? Also, what is the exact complexity of Young-Succinct-Winner and Dodgson-Succinct-Winner?

Acknowledgments: We thank the anonymous referees for their helpful comments and suggestions. This work was supported in part by a National Science Foundation Graduate Research Fellowship under NSF grant no. DGE-1102937.

References

[BD09] N. Betzler and B. Dorn. Towards a dichotomy of finding possible winners in elections based on scoring rules. In *Proceedings of the 34th International Symposium on Mathematical Foundations of Computer Science*, pages 124–136. Springer-Verlag *Lecture Notes in Computer Science #5734*, August 2009.

[Bla58] D. Black. *The Theory of Committees and Elections*. Cambridge University Press, 1958.

[BTT89] J. Bartholdi, III, C. Tovey, and M. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.

[BTT92] J. Bartholdi, III, C. Tovey, and M. Trick. How hard is it to control an election? *Mathematical and Computer Modeling*, 16(8/9):27–40, 1992.

[CSL07] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate? *JACM*, 54(3):Article 14, 2007.

[DKNS01] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar. Rank aggregation methods for the web. In *Proceedings of the 10th International World Wide Web Conference*, pages 613–622. ACM Press, March 2001.

[Dod76] C. Dodgson. A method of taking votes on more than two issues. Pamphlet printed by the Clarendon Press, Oxford, and headed "not yet published" (see the discussions in [MU95, Bla58], both of which reprint this paper), 1876.

[FHH09] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. How hard is bribery in elections? *JAIR*, 35:485–532, 2009.

[FHH15] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. Weighted electoral control. *JAIR*, 52:507–542, 2015.

[FHHR09] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Llull and Copeland voting computationally resist bribery and constructive control. *JAIR*, 35:275–341, 2009.

[FHHR11] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. *Information and Computation*, 209(2):89–107, 2011.

[Gib73] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41(4):587–601, 1973.

[GJ79] M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1979.

[Hem89] L. Hemachandra. The strong exponential hierarchy collapses. *JCSS*, 39(3):299–322, 1989.

[HH07] E. Hemaspaandra and L. Hemaspaandra. Dichotomy for voting systems. *JCSS*, 73(1):73–83, 2007.

[HHR97] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Exact analysis of Dodgson elections: Lewis Carroll's 1876 voting system is complete for parallel access to NP. *JACM*, 44(6):806–825, 1997.

[HHR09] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Hybrid elections broaden complexity-theoretic resistance to control. *Mathematical Logic Quarterly*, 55(4):397–424, 2009.

[HHR14] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. The complexity of online manipulation of sequential elections. *JCSS*, 80(4):697–710, 2014.

[HHS14] E. Hemaspaandra, L. Hemaspaandra, and H. Schnoor. A control dichotomy for pure scoring rules. In *AAAI-14*, pages 712–720, July 2014.

[HS16] E. Hemaspaandra and H. Schnoor. Dichotomy for pure scoring rules under manipulative electoral actions. In *ECAI-16*, pages 1071–1079, Aug/Sep 2016.

[HSV05] E. Hemaspaandra, H. Spakowski, and J. Vogel. The complexity of Kemeny elections. *Theoretical Computer Science*, 349(3):382–391, 2005.

[Kar72] R. Karp. Reducibilities among combinatorial problems. In R. Miller and J. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103, 1972.

[Kem59] J. Kemeny. Mathematics without numbers. *Daedalus*, 88:577–591, 1959.

[Len83] H. Lenstra, Jr. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8(4):538–548, 1983.

[Lin12] A. Lin. *Solving Hard Problems in Election Systems*. PhD thesis, Rochester Institute of Technology, Rochester, NY, 2012.

[McG53] D. McGarvey. A theorem on the construction of voting paradoxes. *Econometrica*, 21(4):608–610, 1953.

[MU95] I. McLean and A. Urken. *Classics of Social Choice*. University of Michigan Press, 1995.

[MW13] N. Mattei and T. Walsh. PrefLib: A library for preferences. In *ADT-13*, pages 259–270, 2013.

[RSV03] J. Rothe, H. Spakowski, and J. Vogel. Exact complexity of the winner problem for Young elections. *Theory of Computing Systems*, 36(4):375–386, 2003.

[Sat75] M. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.

[Wag87] K. Wagner. More complicated questions about maxima and minima, and some closures of NP. *Theoretical Computer Science*, 51(1–2):53–80, 1987.

[YL78] H. Young and A. Levenglick. A consistent extension of Condorcet's election principle. *SIAM Journal on Applied Mathematics*, 35(2):285–300, 1978.

[You77] H. Young. Extending Condorcet's rule. *Journal of Economic Theory*, 16(2):335–353, 1977.

A Appendix

Consider the following election system X. Every voter votes by an approval vector (i.e., states an approval or disapproval for each candidate). A candidate wins if they are approved by exactly half of the voters.

Theorem 19 *X-CCAC is in P and X-Succinct-CCAC is NP-complete.*

Proof. This problem is basically subset sum, which is in P for unary numbers and NP-complete for binary numbers [GJ79]. To show that X-CCAC is in P, let $s_1, \ldots, s_m$ be the approval scores of candidates $c_1, \ldots, c_m$, let the c_1 be the preferred candidate and let $c_1, \ldots, c_\ell$ be the registered candidates. We want to know if there exists a subset I of $\{\ell+1, \ldots, m\}$ of size at most k such that

$$2s_1 = \sum_{1 \le i \le \ell} s_i + \sum_{i \in I} s_i.$$

This can easily be computed in polynomial time, using dynamic programming. Simply let, for all $\ell+1 \le i \le m$, $0 \le k' \le k$, $0 \le t \le \sum_{1 \le i \le m} s_i$, $S[i,t,k'] = 1$ if and only if there exists a size-k' subset I of $\{\ell+1, \ldots, m\}$ such that $\sum_{1 \le i \le \ell} s_i + \sum_{i \in I} s_i = t$. This can be done in polynomial time, since there are only polynomially many i, k', and t to consider.

When the input is given succinctly, there are exponentially many values for t possible. In this case, we can prove NP-completeness. Membership in NP is immediate. We reduce from Partition: Given m positive integers $k_1, \ldots, k_m$ that sum to $2K$, we need to determine if a subset of these integers sums to K. We will have $m+1$ candidates, preferred candidate p and candidates $c_1, \ldots, c_m$. We have K voters that approve of only p, and k_i voters that approve of only c_i. p is the only registered voter and the addition limit is m. It is immediate that there is a subset of $k_1, \ldots, k_m$ that sums to K if and only if p can be made a winner by adding candidates. ❑

A.1 Proof of Theorem 6

Proof. Given a set of candidates C, a list of registered votes V, a list of unregistered votes U, an addition limit $k \in \mathbb{N}$, and a preferred candidate $p \in C$ we can determine if the chair can ensure that p wins in polynomial time. The argument closely follows the approach used by Faliszewski et al. [FHH09] to show that m-candidate α-bribery and m-candidate α-CUCM, for every scoring rule α are each in P for succinct elections.

We describe an integer linear program with a fixed number of variables, which due to the result from Lenstra Jr. [Len83], can be solved in polynomial time.

Since the number of candidates m is fixed, we know that only $m!$ different preference orders are possible. We list these preference orders in order in the following way: $o_1, \ldots, o_{m!}$. For $1 \le i \le m!$, let k_i be the number of voters in V with preference order o_i, and let k'_i be the number of voters in U with preference order o_i. Let $x_1, \ldots, x_{m!}$ be variables. The

constants used as input to our linear program are the coefficients of the scoring vector $\alpha_1, \ldots, \alpha_m$, the counts for the registered voters $k_1, \ldots, k_{m!}$, the counts for the unregistered voters $k'_1, \ldots, k'_{m!}$, and the addition limit k. We have the following constraints.

We first need to ensure that all of our variables have a positive value. So, $\forall i, 1 \leq i \leq m_!$,

$$x_i \geq 0.$$

We also need to ensure that no more than the number of unregistered votes of each type are added. So, $\forall i, 1 \leq i \leq m!$,

$$x_i \leq k'_i.$$

The following constraint ensures that no more than k total votes are added.

$$\sum_{i=1}^{m!} x_i \leq k.$$

Our final constraint ensures that no other candidate has a score greater than p. (Note that $\text{pos}(c, j)$ below denotes the position of candidate c in preference order o_j.) So, $\forall i, 1 \leq i \leq m$,

$$\sum_{j=1}^{m!} \alpha_{\text{pos}(p,j)}(k_j + x_j) \geq \sum_{j=1}^{m!} \alpha_{\text{pos}(c_i,j)}(k_j + x_j).$$

It is easy to see that the above constraints are satisfied if and only if it is possible to add at most k unregistered voters from U such that p wins.

The above integer linear program can be easily modified for the case of deleting voters. ❑

A.2 Proof of Lemma 9: Weighted-Min-Card-Vertex-Cover-Compare polynomial-time many-one reduces to Kemeny-Succinct-Winner.

As mentioned in the core text, this follows by careful inspection and modification of the proof from [HSV05] that Min-Card-Vertex-Cover-Compare polynomial-time many-one reduces to Kemeny-Winner. That proof consists of a chain of three reductions:

1. Min-Card-Vertex-Cover-Compare reduces to Vertex-Cover-Member.
2. Vertex-Cover-Member reduces to Feedback-Arc-Set-Member.
3. Feedback-Arc-Set-Member reduces to Kemeny-Winner.

We define the following weighted versions of the two intermediate problems.

Name: Weighted-Vertex-Cover-Member

Given: Vertex-weighted graph G and vertex v in G.

Question: Is v a member of a minimum-weight vertex cover of G?

Name: Weighted-Feedback-Arc-Set-Member

Given: Irreflexive and antisymmetric edge-weighted digraph G and vertex v in G.

Question: Is there a minimum-weight feedback arc set of G that contains all arcs entering v?

And we show the following.

1. Weighted-Min-Card-Vertex-Cover-Compare reduces to Weighted-Vertex-Cover-Member (Lemma 20).
2. Weighted-Vertex-Cover-Member reduces to Weighted-Feedback-Arc-Set-Member (Lemma 21).
3. Weighted-Feedback-Arc-Set-Member reduces to Kemeny-Succinct-Winner (Lemma 22).

Lemma 20 *Weighted-Min-Card-Vertex-Cover-Compare polynomial-time many-one reduces to Weighted-Vertex-Cover-Member.*

Proof. We extend the construction from Lemma 4.12 from [HSV05] to vertex-weighted graphs. Given two vertex-weighted graphs G and H with $\omega(G) = \omega(H)$, construct graph F as follows. $F = (G + (\{v\}, \emptyset)) \times (H + (\{w\}, \emptyset))$, keeping the weights of the vertices in G and H, and letting $\omega(v) = \omega(w) = 1$. Map (G, H) to (F, w). Write mvc for the minimum vertex cover weight of a vertex-weighted graph. We extend the argument from Lemma 4.12 from [HSV05] to show that this mapping reduces Weighted-Min-Card-Vertex-Cover-Compare to Weighted-Vertex-Cover-Member.

Note that V is a vertex cover of F if and only if

1. V contains all vertices in $G + (\{v\}, \emptyset)$ and a vertex cover of H. The lightest such vertex covers have weight $\omega(G) + 1 + mvc(H)$ and do not contain w.
2. V contains all vertices in $H + (\{w\}, \emptyset)$ and a vertex cover of G. The lightest such vertex covers have weight $mvc(G) + \omega(H) + 1$ and contain w.

Since $\omega(G) = \omega(H)$ it is immediate that $mvc(G) \leq mvc(H)$ if and only if there is a minimum-weight vertex cover of F that contains w. ❑

Lemma 21 *Weighted-Vertex-Cover-Member polynomial-time many-one reduces to Weighted-Feedback-Arc-Set-Member.*

Proof. We extend the construction from the proof of Lemma 4.8 from [HSV05], which itself is an extension of the standard reduction from Vertex-Cover to Feedback-Arc-Set from [Kar72], to graphs with weights. Given a vertex-weighted graph G, define edge-weighted digraph $H = (W, A)$ as follows.

1. $W = \{v, v' \mid v \in V(G)\}$.
2. A consists of the following edges.
   (a) For $v \in V(G)$, edge (v, v') with weight $\omega(v)$.
   (b) For $\{v, w\} \in E(G)$, edges (v', w) and (w', v) with weight $\omega(G)$.

The reduction maps $(G, \hat{v})$ to $(H, \hat{v}')$.

If V' is a vertex cover of G that contains $\hat{v}$, then, by the proof of Lemma 4.9 from [HSV05], $\{(v, v') \mid v \in V'\}$ is a feedback arc set for H that contains all arcs entering $\hat{v}'$. Note that the weight of this feedback arc set is $\omega(V')$.

Now suppose that A' is a feedback arc set of H that contains $(\hat{v}, \hat{v}')$. Again by the proof of Lemma 4.9 from [HSV05], $V' = \{v \in V \mid \exists w \in W : (v, w) \in A' \text{ or } (v', w) \in A'\}$ forms a vertex cover of G that contains $\hat{v}$. Note that $\omega(V') \leq \omega(A')$. ❑

Lemma 22 *Weighted-Feedback-Arc-Set-Member polynomial-time many-one reduces to Kemeny-Succinct-Winner.*

Proof. Given an irreflexive and antisymmetric edge-weighted digraph G and vertex v in G, we first multiply all edge weights by 2 (this does not change the membership of (G, v) in Weighted-Feedback-Arc-Set-Member). We then use McGarvey's construction [McG53] to construct in polynomial time a succinct election that has G as its weighted majority graph. The remainder of the proof is identical to the proof of Lemma 4.2 of [HSV05]. ❑

A.3 Proof of Theorem 12: Kemeny-Winner for four voters is Θ_2^p-complete.

We combine the constructions from [HSV05] and [DKNS01], to reduce Feedback-Arc-Set-Member to Kemeny-Winner for four voters.

Given an irreflexive and antisymmetric digraph $G = (V, A)$ and vertex $v \in V$, we first compute graph G' as done in [DKNS01]. $G' = (V', A')$ such that $V' = V \cup A$ and $A' = \{(v, (v, w)), ((v, w), w) \mid (v, w) \in A\}$. Note that G has a feedback arc set of size k if and only if G' has a feedback arc set of size k. In addition, and important for our proof, (G, v) is in Feedback-Arc-Set-Member if and only if (G', v) is in Feedback-Arc-Set-Member. Also note that G' is irreflexive and antisymmetric. Now apply the reduction from Feedback-Arc-Set-Member to Kemeny-Winner from [HSV05] on (G', v). This reduction outputs an election $g(G')$ that corresponds to G' and (G', v) is in Feedback-Arc-Set-Member if and only if $(g(G'), v)$ is in Kemeny-Winner. The construction from [DKNS01] shows that we need only four voters in $g(G')$.

A.4 Proof of Theorem 13: Kemeny-Winner for two voters who vote a partial order and Kemeny-Winner for one voter who votes a binary relation are Θ_2^p-complete.

The definition of Kemeny winner can naturally be extended to voters voting any binary relation, in the following way.

A candidate p is a Kemeny winner if p is ranked first in a Kemeny consensus. A Kemeny consensus is a linear order $>$ over the candidates that minimizes the Kendall's tau distance to V, i.e., that minimizes

$$\sum_{a,b \in C, a>b} \|\{v \in V \mid b >_v a\}\| + \|\{v \in V \mid b \sim_v a\}\|/2,$$

where $>_v$ is the preference relation of voter v.

The proof of Θ_2^p-completeness proceeds as in the proof of Theorem 12. All that is needed is to observe that

1. $g(G')$ needs only two voters v and w if voters can vote with a partial order: $x >_v (x, y)$ for all $(x, y) \in A$, $(x, y) >_w y$ for all $(x, y) \in A$, and v and w are indifferent for all other pairs of candidates.

2. $g(G')$ needs only one voter v if voters can voter any binary relation: Let $x >_v y$ if and only if $(x, y) \in A'$.